

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:53:54

PAGE 1

REFERENCE NO: 240

This contribution was submitted to the National Science Foundation as part of the NSF CI 2030 planning activity through an NSF Request for Information, https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf17031. Consideration of this contribution in NSF's planning process and any NSF-provided public accessibility of this document does not constitute approval of the content by NSF or the US Government. The opinions and views expressed herein are those of the author(s) and do not necessarily reflect those of the NSF or the US Government. The content of this submission is protected by the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>).

Author Names & Affiliations

- Jean-Francois Lamarque - National Center for Atmospheric Research, Boulder, CO
- Richard Loft - National Center for Atmospheric Research, Boulder, CO
- Mariana Vertenstein - National Center for Atmospheric Research, Boulder, CO
- Leo Donner - Geophysical Fluid Dynamics Laboratory, Princeton, NJ

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Climate

Title of Submission

High-resolution Earth System Modeling

Abstract (maximum ~200 words).

In this RFI response the National Center for Atmospheric Research focuses on the computational and data requirements and challenges faced by the Climate Modeling Enterprise (CME). Earth System Models (ESMs) are complex multiscale multi-physics software systems developed over decades through the collaboration of many institutions. With low floating point intensity and the need for high throughput rates, as measured by simulated years per day, ESMs are problematic for current post-Dennard-scaling computer architecture. The CME is data intensive, especially on the model output side, and this creates a significant analysis performance issue for extant systems.

Managing the rapidly changing and complex ESM software is itself a grand challenge. ESM developers face requirements for modularity, maintainability, bit-for-bit reproducibility, performance portability, and scaling from laptops to the largest supercomputers available.

To address these needs, we argue that part of NSF's response to NSCI should be to develop purpose-built data-centric systems that integrate simulation and data analysis. An overall integration of the data infrastructure of NSF supercomputing facilities would likely prove fruitful in fostering advances in these important science problems. We close with a discussion of the workforce development, diversity and inclusion issues facing not just the CME, but all of computational science.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

In order to make simulation results from Earth System Models (ESMs) most societally-relevant, simulations need to be performed at resolutions that explicitly capture the processes and phenomena of interest. Increasingly, the scales of interest are reaching within the realm of regional modeling, with the expectation that such models can accurately capture the behavior of weather systems. Under this constraint, we have identified two research challenges associated with high-resolution ESMs

- 1) Building and running high-resolution ESMs, including a documentation of the full simulation uncertainty
- 2) Generating new scientific insight from high-resolution model output analysis

Earth-System Models tend to be developed on a 5- to 7-year cycle, bringing together model developments, increases in model resolution (horizontal and/or vertical, including vertical extent) and/or improved numerical techniques to increase model fidelity, capability and usability. Because of the increasing complexity of the physical/biological/chemical couplings captured in ESMs, the process of building an ESM requires the understanding of the behavior of the model as a coupled system, limiting the utility of working on each component (i.e., atmosphere, ocean, land, sea-ice and land-ice) in isolation. As a consequence, in order to understand the behavior of the most recent version of the Community Earth System Model (CESM2), it has been necessary to perform fully-coupled simulations over several thousand years in order to define a viable configuration for all components of the earth system. This was possible only because a throughput of 20 simulated years per wall-clock day (synd) was obtained on the NSF Yellowstone supercomputer for a ~100km resolution configuration of CESM2. Such throughput is currently unavailable on any supercomputer for the model resolutions required to explicitly simulate important processes, such as those associated with oceanic eddies, or high-impact events such as tropical cyclones. In both cases, the horizontal resolution needs to be an order of magnitude finer than in the standard configuration of CESM2.

Modeling groups usually strive to put together a small number of model configurations, in part because of the complex and lengthy process described above. Nonetheless, models such as ESMs are never perfect and it is critical to place generated results within the broad knowledge of uncertainties and biases. It is therefore only through a quantitative and comprehensive exploration of uncertainty that progress will be made in our understanding of the limitations of current state-of-the-art ESMs. The main sources of limitation are:

- Horizontal resolutions will remain too coarse to capture many of the key processes in the climate system: for instance, it has been shown that a resolution of at least 25-km in the atmosphere model is needed to simulate the intensity, track and statistics of tropical storms, but current climate models generally operate in the 100-km range. Even at 25 km, however, cloud scale processes will still have to be parameterized. Consequently, in the foreseeable future, improving resolution will only incrementally improve our explicit representation of processes and thus, our ability to simulate high-impact events like tropical cyclones. Similar examples in other ESM components include ocean mesoscale eddies or plant ecosystem dynamics.
- Internal chaotic variability is under-sampled: ensemble sizes are too small. With a few exceptions, most notably the CESM Large Ensemble, most published studies are based on 1-10 ensemble members. The ensemble sizes needed may be as much as an order of magnitude larger in order to fully capture the statistical range of climate noise over the whole range of processes and variables of interest. Analysis of internal variability becomes an even greater issue with higher resolution as the number of degrees of freedom increases.
- Parametric uncertainty is highly under-sampled: climate modeling centers typically deliver a model configured for a single set of parameter choices, whereas the space of permitted parameter choices within the bounds of observational uncertainty is much larger. Moreover, some empirical parameters input to models are poorly constrained by data. Systematic explorations of this uncertainty space could yield broad probability distributions of outcomes, but would lead to an explosion in computational cost and data volume.
- Estimates of knowledge uncertainty (also called structural uncertainty) are provided through multi-model analysis, which only provides feedback to modeling centers over a slow cycle of largely uncoordinated analysis and publication by multiple authors. Some modeling centers have also performed identical simulations with two (or more) configurations that differ only from a single component, such as the various ocean models used by the Geophysical Fluid Dynamics Laboratory for multi-model intercomparison submissions.

Combined together, trying to tackle the uncertainty issues highlighted above represents an enormous and currently unrealizable computational challenge. The issue of parameter choice and optimization (also known as model tuning, see Hourdin et al., 2016; Schmidt et al., 2017) becomes even more problematic when models (especially atmosphere or ocean) are using resolutions that are straining current or near-term computational capabilities. It is our firm conviction that only by fully addressing the uncertainty quantification will we be able to accelerate progress on assessing the robustness of ESM projections and predictions, along with defining a path towards model improvements.

The second research challenge relates to the use of high-resolution ESM generated output for scientific discovery. The Climate Model

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:53:54

PAGE 3

REFERENCE NO: 240

Intercomparison Project (CMIP, currently in its sixth phase, see <https://www.wcrp-climate.org/wgcm-cmip/wgcm-cmip6> for more details) provides the research community with a set of simulations that: 1) tackle scientific issues of interest, 2) provide a well-defined protocol to ensure simulations from different modeling groups can be directly compared, and 3) specify a common output format. Such wide-spread effort has been critical in providing the research community (at universities and federal laboratories) with the necessary data to answer many important climate science questions. The key element is a widespread distribution of, and easy access to, the generated datasets so that the research community can perform its analysis. Under the CMIP5 (CMIP Phase 5) protocol, approximately 5 petabytes (PB) of data were generated. While this was a challenge for the modeling centers, it is clear that few analysis groups, usually researchers in a university department, have the resources to download a significant portion of this archive. Indeed, such high volume leads to computational bottlenecks for the analysis of generated model output. These bottlenecks consist of: 1) access to sufficient disk space to store data to be analyzed by university researchers, and 2) access to sufficient computational power to perform timely scientific inquiries of the data. While many high-quality works have been published under these restrictions, it is also true that other important questions cannot be addressed. Furthermore, with increasing complexity (more output fields) and more emphasis on output at high-resolution in time and space to capture extreme events, the bottlenecks will become even more pronounced. This will strongly reduce the number and breadth of the analyses of ESM output simulations needed to answer the societally-relevant questions.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

The nonlinear and chaotic nature of the Earth's climate system, combined with its intrinsic complexity, severely limits the use of analytical techniques to gain insight into its behavior: for this reason climate science and computer modeling have become inseparably linked. The design, methodology, and capabilities of ESMs are largely governed by the nature and capabilities of the available compute environment. Thus it is of paramount importance to sustaining scientific progress that the software and the computing environment harmoniously coevolve.

Disruptive trends computing architecture, prompted by the end of Dennard scaling and the advent of multi-core and many-core processors with Single Instruction Multiple Data (SIMD) features over the last decade, have increased the need to continuously refactor existing codes for optimal performance and scalability. With rapidly changing model codes and limited developer resources, modelers are unwilling to support multiple versions of software for different architectures, and this further constrains optimization approaches. Even so, single source code-level refactoring efforts have yielded gains of 2x-4x, and many of these optimizations have been incorporated into production models. Others have not, either for readability or correctness reasons, or because in some cases these optimizations have exhibited reduced scalability on the thread and MPI process level, reducing or completely nullifying the benefits of core-level optimization. While any improvements are welcome, it is worth noting that these results fall far short of the orders of magnitude improvements required to advance the CME in the coming years.

The current trends in computing architectures best support weakly-scalable workloads with high floating point intensity (FPI) - the opposite of what is required by the CME's Earth System Models, where throughput, as measured by simulated years per day (SYPD), is the relevant metric of scientific throughput. To achieve high SYPD, ESMs with large numbers of state variables, load imbalances, and low FPIs are run under strong scaling conditions. Thus ESMs are most sensitive to memory bandwidth and interconnect latency. HPC architectures design trends that further add flops/byte are unhelpful to low computational intensity codes. Most of the performance advantages currently seen in GPUs and Intel Xeon Phi systems running meteorological workloads come from the higher memory bandwidth of such processors, either from GDDR or more recently from stacked, high bandwidth memory (HBM) deployed with them.

The computational bottleneck the CME are just now facing is deeply rooted in the underlying device physics and the techno-economics of the computing marketplace and, in our view, is unlikely to be strongly influenced by one agency's, e.g. NSF, or even one government's efforts to stimulate a workaround through co-design efforts. A broader initiative with a very large incentive might help vendors address the techno-economics, but may be not be able to overcome the limitations of underlying device physics.

A more tractable approach might be to increase funding of algorithmic innovations - for example, more effective time integration techniques, replacing model parameterizations with machine learning encoders or, alternatively, by exploring reduced precision physics as Tim Palmer of Oxford University and others have suggested.

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:53:54

PAGE 4

REFERENCE NO: 240

Because of the long road to adoption of new methods, both research and model development funding must be provided to support the deployment of algorithmic improvements in production applications. Otherwise, the end result will be simply adding to the steadily growing and conflicting software development to-do list of already overloaded ESM development teams with fixed (or declining) resources.

Over the last 15 years the complexity of ESMs, as measured by source lines of code (SLOC), has grown by an order of magnitude. As the size and complexity of Earth Systems Models continues to increase, the need to be able to maximize code reuse between different modeling efforts is becoming paramount. However, today, many model components and sub-components are tied into their host model's infrastructure, preventing code reuse. Using a non-interoperable component in a new host model requires re-implementation of the parts of the component which interacted with the old host model's infrastructure. This process is both time consuming and error prone. Interoperable infrastructure allows model components (e.g., an atmosphere model) and sub-components (e.g., a radiation parameterization) to easily be moved between models. Thus interoperability allows code reuse while leveraging focused infrastructure improvements across participating models.

Much of the increasingly complex functionality of today's models is expressed in code logic. However, this logic increases the number of possible code paths geometrically, a trend that is unsustainable. Finally, any modifications to code functionality (e.g., parameterization sequencing) is difficult, error prone, and results in a new model to maintain. Re-implementing model functionality with data-driven scheduling infrastructure allows model complexity and scheduling to be captured in a simple data file which also provides documentation of any experiment's code path. Furthermore, since every experiment type is captured in a simple data file, complexity grows linearly with model functionality.

Bit-for-bit (BFB) reproducibility has long been a requirement for climate model development. BFB reproducibility allows testing for software quality (e.g., use of uninitialized variables, errors in parallel decomposition, race conditions) and also ensures that model simulations can be continued from a previously-saved state (i.e., model restart). Sometimes, BFB reproducibility is not possible, for instance, some important code optimizations or parameterization updates which are not expected to change the simulated climate. In these cases, specialized tools are required to distinguish code modifications which change the simulated climate from answer changes which still represent the same climate.

Coupled-climate model development is a lengthy, complex, and costly process. To facilitate different stages of development, the model configurations must run on everything from a laptop (e.g., simplified model version used for parameterization development) to the largest computational platforms available (e.g., for multi-century fully-coupled ensemble climate experiments). The computational cost of the model-development process is enormous, placing a premium on code portability and reusability. It is simply unfeasible to attempt to re-engineer the code for each new computer architecture. Therefore, any change in code structure must be carefully considered to ensure long-term viability.

Data assimilation is an essential capability in the drive to use coupled-climate models for seasonal, subseasonal forecasts and decadal prediction experiments. This capability, however, adds new heavy I/O loads and new computational burdens on climate simulations. These needs must be taken into account when planning both software infrastructure improvements (e.g., streamlining I/O for data assimilation) and hardware requirements.

The CME's analysis workflow involves the study of massive volumes of output data. For example, NCAR's Cheyenne system is expected to produce 2-3 petabytes/month, much of it climate output. This torrent of data is typically averaged, re-gridded, reshuffled into time series, recombined into derived quantities, and may be subset. In the end, the goal of all is to gain new scientific insights into the workings of the Earth System, uncover model deficiencies, and subsequently design model improvements.

The increase in ESM data output volumes and model complexity have outstripped the capabilities of data subsystems based on hard disk drives (HDD) and traditional serial monolithic analysis scripts. The introduction of storage systems based on solid state NVMe devices, the growth of the Python software ecosystem, the introduction of platform as a service via containers have created opportunities to dramatically refactor these workflows for better performance.

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:53:54

PAGE 5

REFERENCE NO: 240

We close with a few additional remarks regarding three topics: 1) the need for customizing system architectures for science and engineering applications with similar computational and data analysis requirements; 2) workforce development issues surrounding scientific supercomputing; and 3) the importance of diversity initiatives in computational science as a way of combating a growing “digital divide” and lack of inclusion in computational science fields.

Strawman NSF NSCI Strategy: The case for purpose-built architecture It is likely that the cyberinfrastructure requirements of the CME will be found to be common to many other problems in the physical sciences (e.g. stellar astrophysics, seismology, combustion, aeronautical engineering) in which complex, multi-scale, multi-physics components interact in a nonlinear, coupled system. Therefore, we recommend that NSF focus a sizeable fraction of its HPC CI investments on one or more systems designed to support these important use cases by fostering the development of an integrated simulation-analysis HPC architecture, i.e. one that focuses on providing: 1) a large, stable simulation system composed of multi-core processors; 2) a tightly coupled analysis system; and 3) a tiered storage system. The latter would be, the heart of the system, with an integrated hot tier of storage, likely composed of novel low-latency storage technologies, accessible to both system components of the simulation-analysis complex forming its “scientific scratch pad”. The use of traditional tiers consisting of POSIX file systems running on rotating disk and tape archives should be reconsidered, factoring in new scalable, flat object stores from the cloud as archiving alternatives. The low latency, high bandwidth “hot-tier” of storage leverages truly transformative technology trends and is critical to accelerating data intensive scientific research, such as that found in the CME, as it will enable the rapid multivariate evaluation of large-scale model output both offline and as they are running. Finally, the analysis system should present a virtualized interface to the end user, allowing the easy creation, sharing and documentation of data-centric workflows. An overall integration of the data infrastructure of NSF supercomputing facilities would likely also prove fruitful in fostering advances in these important science problems.

We further recommend that the NSF support the exploration novel statistical, machine learning, and deep learning applications in conjunction with this strategy, so that these new approaches can be evaluated, and where appropriate, incorporated into workflows on the integrated platforms. This initiative would have the objectives of automating certain routine applications of expert scientific knowledge, thus freeing scientists to focus on creative activities, such as extracting and sharing new scientific knowledge.

Workforce Development Requirements It has always been hard to find suitable job candidates with the unique combination of mathematical, scientific and computational skills to contribute to the CME. Factors often cited as barriers to attracting and retaining new talent include compensation, the market value and translatability of the skills acquired (consider how experience in Java and Python vs Fortran plays on a resume). Today, because of large-scale forces at play in the IT industry, it is becoming increasingly difficult to attract software engineers and computational scientists to the scientific computing and data analytics problems found in the CME. Historically, hiring managers in high-performance scientific computing have been able to compensate for the poorer salaries on offer because the work is more intellectually rewarding, societally impactful, and challenging. Today, this intellectual high ground can arguably also be claimed by the commercial space. The Googles, Amazons and Apples can also attract people with challenging and societally impactful problems. As a result, it is becoming increasingly difficult to recruit new talent in computational science fields like data science, GPU programming, machine learning and deep learning.

Enhancing Diversity and Inclusion The CME require a diversity of perspectives, approaches and viewpoints to succeed. However, in both the geosciences domain and computer science, mathematics and related disciplines entrained by the CME, the workforce continues to be plagued by consistently low, and in some cases declining, proportions of women. According to the NSF’s latest (2017) “Women, Minorities, and Persons with Disabilities in Science and Engineering” report, the proportion of women receiving Ph.D.s is stuck at 10%, roughly the same as in 1994. Women make up just 8.3% of the computer science faculty at universities. Diversity figures are not much better in industry, hovering at around 15% of technical staff.

In recent years, Silicon Valley firms have been addressing these issues via company-wide implicit bias training and large education and outreach investments (in 2015 Intel pledged \$300M over five years), and technological innovations in the hiring process designed to obscure a candidate’s gender and race. These efforts seem to be moving the needle in the private sector and a similar approach, guided by the latest diversity research, may be successful in academia and the sciences. It is clear from the lack of progress in academia that a new approach and nontrivial investments in diversity are required to effect change. While the NSF should be praised for its development of the INCLUDES program, more investments, perhaps in partnership with the leading universities and laboratories, may be required to crack the diversity barriers at leading Computer Science and Geoscience departments.

Another research study in the journal Science found that the workplace representation of women and minorities by discipline was best predicted by the degree to which that discipline believed in “innate talent” or “genius” as the predictor of success in it. Thus, since some professions, that viewed themselves as “trainable”, (e.g. the law) have better diversity scores, whereas ones like physics, computer science

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:53:54

PAGE 6

REFERENCE NO: 240

and math, where innate talent or genius is viewed as key, diversity is low. The study suggests that disciplinary self image may be a contributing factor in implicit bias, and could form the basis for implicit bias training in an academic setting.

Finally, the NSF should focus on programs for supercomputing and related disciplines that go far beyond head counts, to ones that track and encourage the long term inclusion of diverse groups in the workplace. These efforts should include both qualitative and quantitative metrics of diversity over time measured for various types of diversity programs and institutions, providing a means of distinguishing effective programs from ineffective ones and superficial from substantive institutional change.

Consent Statement

- "I hereby agree to give the National Science Foundation (NSF) the right to use this information for the purposes stated above and to display it on a publically available website, consistent with the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>)."
-